

Description

METHOD AND APPARATUS FOR PROVIDING FAULT TOLERANCE TO MEMORY

BACKGROUND OF INVENTION

[0001] 1. Field of the Invention

[0002] The invention relates to a method and apparatus for providing fault tolerance to a memory and more particularly, to a method and a networking apparatus, such as a switch, a router, or the like, for providing fault tolerance to a memory.

[0003] 2. Description of the Prior Art

[0004] A networking apparatus such as a switch, a router, or the like, is usually employed in the construction of a network of computers. One of the functions of the networking apparatus is to efficiently forward data sent by users (computers on the network) to the data's destination via the ports of the networking apparatus. To accomplish this

task, information about host/port relationships must be stored in the memory of the networking apparatus, which is often in the form of a look-up table called a MAC address table.

[0005] The MAC address table records the source MAC IDs (SIDs) of incoming packets. If there is an incoming packet to be stored in the MAC memory, a rule called "hashing" is defined for mapping the incoming packet to a specific look-up table entry based on the MAC ID of the incoming packet. Whenever a packet enters the networking apparatus, the networking apparatus picks the destination MAC ID (DID) of this incoming packet and uses DID to perform hashing. Hashing of the DID maps this DID to a specific entry in the look-up table. The networking apparatus checks the MAC ID stored in this entry to see whether this MAC ID appeared before and has been recorded in the table, this step being called "search". If the MAC ID stored in this entry is invalid, out of date, or valid but different from the DID under searching, the search result is "missed". However, if the MAC ID stored in this entry is valid and indeed equal to the DID of the incoming packet, the search result is "hit". The networking apparatus decides how to forward the packet based on the search result. If the

search result is "missed", and this packet is a legal packet, this incoming packet will be broadcasted to all the ports except the inbound port. If the search result is "hit", this incoming packet will be forwarded to the associated port. This step is called "forward". After forwarding, the network apparatus would try to "learn" the "SID-inbound port" relationship and record the relationship into the look-up table.

[0006] One of the aims of utilizing the look-up table in the MAC memory is to prevent unnecessary broadcasting of the incoming packets so that the amount of data transmission can be reduced. Because of the important role the MAC memory storing the host/port relationship plays, current networking apparatus employ only memory free of defects.

[0007] However, the fabrication process of memory is by no means perfect. In other words, the yield of memory with no physical defects from the manufacturing process of memory is not 100%. For every batch of memory fabricated, a substantial portion will prove to be defective and as a result be discarded. The end result is a higher manufacturing cost.

[0008] As for the details concerning hashing scheme, please refer

to "a comparison of hashing schemes for address lookup in computer networks", by Jain, IEEE Transactions on Communications, COM40 (10): 1570–1573, October 1992, which is incorporated herein for reference.

SUMMARY OF INVENTION

- [0009] Therefore, one of the many objectives of the claimed invention is to provide a method and apparatus for providing fault tolerance to a memory.
- [0010] According to the claimed invention, a networking apparatus is disclosed. The networking apparatus for providing fault tolerance to memory comprises a first memory for storing information regarding host/port relationships; and a second memory for indicating the status of the first memory; wherein the networking apparatus marks the second memory based upon the status of the first memory.
- [0011] According to the claimed invention, a method for providing fault tolerance to memory in a networking apparatus is also disclosed. The method comprises performing a built-in self test (BIST) on a first memory when the networking apparatus powers on; marking a second memory to indicate which sections of the first memory is defective; locating one of the sections of the first memory according

to the packet; and checking the portion of the second memory corresponding to the located section of the first memory to determine how to handle the packet.

[0012] One of the many advantages of the claimed invention is the ability of the claimed invention to use defective memory. This ability allows manufacturers to manufacture more networking apparatuses for every batch of memory fabricated. Because of the resulting efficiency, production costs can be lowered.

[0013] These and other objectives of the claimed invention will no doubt become obvious to those of ordinary skill in the art after reading the following detailed description of the preferred embodiment that is illustrated in the various figures and drawings.

BRIEF DESCRIPTION OF DRAWINGS

[0014] Fig.1 is a diagram of a networking apparatus according to one embodiment of the present invention.

[0015] Fig.2 is a flowchart of a process employed by a networking apparatus when handling a packet according to one embodiment of the present invention.

[0016] Fig.3 is a flowchart of a packet transmission process in a networking apparatus when there is no substitute memory present according to one embodiment of the present in-

vention.

[0017] Fig.4 is a flowchart of a packet transmission process in a networking apparatus when there is a substitute memory present according to one embodiment of the present invention.

[0018] Fig.5 is a flowchart of packet information learning process in a networking apparatus when there is no substitute memory present according to one embodiment of the present invention.

[0019] Fig.6 is a flowchart of packet information learning process in a networking apparatus when there is a substitute memory present according to one embodiment of the present invention.

DETAILED DESCRIPTION

[0020] Please refer to Fig.1. Fig.1 is a simplified diagram of a networking apparatus 10 according to an embodiment of the present invention. In this preferred embodiment, the networking apparatus 10 comprises a MAC address memory 20, a status record memory 30, a substitute memory 40, and a processor 50 coupled to the memories 20, 30, 40. The MAC address memory 20 is for storing host/port (SID-inbound port) relationships as used in a conventional networking apparatus, the status record memory 30 is for

indicating the status of the MAC address memory 20, the substitute memory 40 is for serving as a secondary site for storing information regarding host/port relationships, and the processor 50 is for manipulating the memories 20, 30, 40.

[0021] The networking apparatus 10 can be embodied by a switch, a router, or the like. In this embodiment, a switch is taken as an example for the network apparatus 10. The MAC address memory can be embodied by but not limited to SRAM memory. The status record memory 30 can be embodied by but not limited to a group of registers. In this embodiment, each register corresponds to one of the entries of the MAC address memory 20. However, it is not the limitation of the present invention. The substitute memory 40 can be embodied by but not limited to a CAM content addressable memory. The embodiments of these parts are merely to serve as examples and are not meant to act as limitations. The specification will refer to the parts of the invention by their given examples.

[0022] Please note that the substitute memory 40 in this embodiment is an optional part. The following descriptions will detail how to enable the invention with or without the substitute memory 40.

[0023] Please refer to Fig.2. Fig.2 is a flowchart of a base process employed by the networking apparatus 10 when handling a packet according to an embodiment of the present invention. The flowchart shown in fig.2 at least includes the following steps:

[0024] Step 100:Perform a BIST Built In Self-Test. A switch 10 performs a BIST to check if there is any defect present in a SRAM 20 when the switch is first powered on.

[0025] Step 110:Mark Status Record Memory. If any defects are found in the SRAM 20, the switch 10 will mark a corresponding register 30 accordingly for each found defect. The corresponding defects are marked so that the switch 10 will be able to determine if a located entry in the SRAM 20 is suitable for storing information. It should be noted that the marking of the register 30 is not to be limited in a corresponding manner. The mark stored in the registers 30 can be in the form that the specific entry with defect can be identified. Thus, the size and the number of the registers 30 can be reduced.

[0026] Step 120:Locate Entry in MAC Address Memory. Whenever a packet is being handled, the switch 10 performs searching to the packet and locates an entry in the SRAM 20 according to the MAC ID of the incoming packet. In this em-

bodiment, the entry is located based on either the Destination ID or the Source ID of the packet, depending on if the switch is conducting packet transmission or packet information learning. The located entry in the SRAM 20 should indicate which port that the packet should be forwarded to or serve as the entry the packet information should be learned into.

[0027] Step 130:Check Status Record Memory. However, since the located entry in the SRAM 20 may be defective and therefore does not indicate which port that the packet should be forwarded to or cannot serve as the entry the packet information should be learned into, the register 30 is checked. In this embodiment, the switch 10 checks to see how the register 30 corresponding to the located entry is marked. If the marking of the register 30 is not in a corresponding manner, the switch 10 will check if there is at least one of the markings stored in the register 30 identify the specific entry.

[0028] For details concerning BIST, please refer to "a programmable BIST core for embedded DRAM", by Huang et. al., IEEE Design and Test Magazine, Jan-Mar 1999, which is incorporated herein for reference.

[0029] Fig.3 – Fig.6 are a continuation from the process in Fig.2.

The steps that follow Step 130 in Fig.2 depend on what process the switch 10 executes. Please refer to Fig.3. Fig.3 is a flowchart of a packet transmission process in a networking apparatus 10 when there is no substitute memory 40 present according to an embodiment of the present invention.

[0030] Step 130:Check Status Record Memory. However, since the located entry in the SRAM 20 may be defective and therefore, not indicate which port that the packet should be forwarded to, the status record memory is checked. In this embodiment, the switch 10 checks to see how the register 30 corresponding to the located entry is marked. If the register 30 indicates the located entry is defective, go to Step 140. Otherwise, go to Step 150.

[0031] Step 140:Broadcast Packet. Because the register 30 indicates that the located entry is defective or that the content of the located entry does not match, the switch 10 will be unable to find the port that the packet should be forwarded to. Therefore, the switch 10 broadcasts the packet to all ports except for the port that the packet originated from. Go to Step 170.

[0032] Step 150:Compare Packet to MAC Address Memory. Even though the register 30 indicates that the located entry is

not defective, the switch 10 needs to check the content of the located entry of the SRAM 20 by comparing it to the MAC ID of the packet. If the content of the packet and the located entry match, go to Step 160. Otherwise, go to Step 140.

[0033] Step 160:Forward Packet. The content of the located entry of the SRAM 20 matches, and the switch 10 is able to forward the packet to the port in accordance with the content of the located entry of the SRAM 20.

[0034] Step 170:Finish. The handling of the packet is finished, and the switch returns to Step 120 in Fig. 2 to await the handling of the next packet.

[0035] Please refer to Fig.4. Fig.4 is a flowchart of a packet transmission process in a networking apparatus 10 when there is a substitute memory 40 present according to an embodiment of the present invention.

[0036] Step 130:Check Status Record Memory. However, since the located entry in the SRAM 20 may be defective and therefore, not indicate which port that the packet should be forwarded to, the status record memory is checked. If the register 30 indicates the located entry is defective, go to Step 240. Otherwise, go to Step 250.

[0037] Step 240:Search Substitute Memory. Because the register

30 indicates that the located entry is defective, the switch 10 will search the CAM 40 for the information indicating which port that the packet should be forwarded to. In this preferred embodiment, the switch 10 searches the CAM 40 according to the packet for a match with the packet. If the search is successful, go to Step 260. Otherwise, go to Step 270.

[0038] Step 250: Compare Packet to MAC address Memory. Even though the register 30 indicates that the located entry is not defective, the switch 10 needs to check the content of the located entry of the SRAM 20 by comparing it to the packet. If the content of the packet and the located entry match, go to Step 260. Otherwise, go to Step 240.

[0039] Step 260: Forward Packet. The switch 10 forwards the packet to the port in accordance with the content of the located entry of the SRAM 20 or the found entry in the CAM 40. Go to Step 280.

[0040] Step 270: Broadcast Packet. Because the content in the located entry of the SRAM 20 does not match or search in the CAM 40 is unsuccessful, the switch 10 will be unable to find the port that the packet should be forwarded to. Therefore, the switch 10 broadcasts the packet to all ports except for the port that the packet originated from.

[0041] Step 280:Finish. The handling of the packet is finished, and the switch returns to Step 120 in Fig. 2 to await the handling of the next packet.

[0042] Fig.4 differs from Fig.3 in what is done if the corresponding register 30 indicates that the located entry of the SRAM 20 is defective, as well as in what is done if the comparison between the packet and the located entry does not have a match. Instead of directly broadcasting the packet as in Fig.3, Fig.4 has the CAM 40 searched according to the content of the incoming packet. It should be noted that the function of the substitute memory 40 is to provide an additional and substitution space for the MAC address memory 10. If the entry of the MAC address memory 10 is defective, the content of the incoming packet, which ought to be stored in the defective entry can be stored in the substitute memory 40. In this embodiment, the substitute memory 40 is a content addressable memory, CAM. That is, the content of the incoming packet can be stored in the CAM directly if the corresponding entry of the MAC address memory is defective.

[0043] Fig.5 is a flowchart of packet information learning process in a networking apparatus 10 when there is no substitute memory present according to an embodiment of the

present invention.

[0044] Step 130:Check Status Record Memory. However, since the located entry in the SRAM 20 may be defective and therefore does not indicate which port that the packet should be forwarded to, the status record memory is checked. If the register 30 indicates the located entry is not defective, go to Step 340. Otherwise, go to Step 350.

[0045] Step 340:Execute Learning Process In MAC address Memory. Since the register 30 indicates that the located entry in the SRAM 20 is not defective, the switch 10 is able to store information into the located entry. Therefore, the switch 10 executes the learning process in the SRAM 20. The learning process comprises checking the content of the located entry in the SRAM 20, and if the located entry in the SRAM 20 is available, learning the information pertaining to the packet. For example, in this preferred embodiment, the switch 10 checks an ID stored inside the located entry. The switch 10 will learn the SID of the packet in all cases except when the located entry stores an ID, which is valid but different from the Source ID of the packet.

[0046] Step 350:Finish. The handling of the packet is finished, and the switch returns to Step 120 in Fig. 2 to await the

handling of the next packet.

[0047] Fig.6 is a flowchart of packet information learning process in a networking apparatus 10 when there is a substitute memory present according to an embodiment of the present invention.

[0048] Step 130:Check Status Record Memory. However, since the located entry in the SRAM 20 may be defective and therefore does not indicate which port that the packet should be forwarded to, the status record memory is checked. If the register 30 indicates the located entry is not defective, go to Step 440. Otherwise, go to Step 450.

[0049] Step 440:Check Located Entry. Since the register 30 indicates that the located entry in the SRAM 20 is not defective, the switch 10 is able to store information into the located entry. Therefore, the switch 10 will execute the learning process. However, before the learning process, the switch 10 first checks the content stored in the located entry for validity and value. If the stored ID is valid but different from the SID of the packet, go to Step 450; otherwise, go to Step 445.

[0050] Step 445:Execute Learning Process In MAC Address Memory. In this embodiment, when the located entry in the SRAM 20 is available, the switch 10 will learn the SID of

the packet in all cases except when the located entry stores an ID, which is valid but different from the Source ID of the packet. Go to Step 470.

[0051] Step 450:Search Substitute Memory. Following Step 130 or Step 440, the switch 10 will attempt to learn information pertaining to the packet into the CAM 40. In the preferred embodiment, the switch 10 searches the CAM 40 according to the packet. If the located entry of the CAM 40 is available, go to Step 460. Otherwise, go to Step 470.

[0052] Step 460:Learn Information Into Substitute Memory. The switch 10 learns information pertaining to the packet into the located entry of the CAM 40. The information learned is the same as in Step 440.

[0053] Step 470:Finish. The handling of the packet is finished, and the switch returns to Step 120 in Fig. 2 to await the handling of the next packet.

[0054] Fig.6 differs from Fig.5 in what is done if the corresponding register 30 indicates that the located entry of the SRAM 20 is defective as well as in what is done if the stored ID is found to be valid but different from the SID of the packet. Instead of going directly to the finishing step as in Fig.5, the switch 10 searches the CAM 40 according to the content of the packet.

[0055] In other words, Fig.2 Fig.6 illustrate how a packet is handled. To be specific, Fig. 2 and Fig.3 detail packet transmission using only a SRAM 20 and set of registers 30. Fig.2 and Fig.4 also detail packet transmission but with the addition of a CAM 40. Fig.2 and Fig.5 detail packet information learning process using only a SRAM 20 and a set of registers 30. Fig.2 and Fig.6 detail packet information learning process but with the addition of a CAM 40.

[0056] Please note that the reception of a packet can cause the switch 10 to conduct two processes. For example, upon receiving a packet, the switch 10 can execute the packet transmission and then return to the locating step in Fig.2 and execute packet information learning before returning to a finishing step. Also, the order of the two processes can be reversed with packet information learning occurring before packet transmission.

[0057] As one can see, the embodiments of the present invention allow manufacturers to incorporate defective memory into the networking apparatus. Using the information disclosed, the networking apparatus with defective memory is able to function in substantially the same way to achieve substantially the same result as a networking apparatus with defective-free memory.

[0058] Those skilled in the art will readily observe that numerous modifications and alterations of the device may be made while retaining the teachings of the invention. Accordingly, that above disclosure should be construed as limited only by the metes and bounds of the appended claims.